

WHAT IS CLAIMED IS:

1. A method of managing I/O resources in an information delivery environment, comprising
5 modeling utilization of at least one of said I/O resources; and managing at least one of said I/O
resources based at least in part on said modeled utilization.

2. The method of claim 1, wherein said modeling comprises modeling utilization of at least
10 one of said I/O resources based at least in part on at least one of said system I/O performance
characteristics associated with said I/O resources.

3. The method of claim 2, wherein the value of at least one of said system I/O performance
15 characteristics is estimated.

4. The method of claim 2, wherein the value of at least one of said system I/O performance
characteristics is monitored.

5. The method of claim 1, wherein said I/O resources comprise at least one of said file
system resources, storage system resources, or a combination thereof.

6. The method of claim 1, wherein said information delivery environment comprises
25 delivery of continuous media data from an information management system to a network;
wherein said I/O resources comprise I/O capacity and buffer memory space of said information
management system; and wherein said managing comprises balancing said I/O capacity with said
buffer memory space to ensure uninterrupted delivery of said continuous media data.

7. The method of claim 1, wherein said information delivery environment comprises delivery of continuous media data from an information management system to a plurality of viewers across a network; wherein said I/O resources comprise I/O capacity and buffer memory space of said information management system; and wherein said managing comprises balancing said I/O capacity with said buffer memory space to ensure uninterrupted delivery of said continuous media data to said plurality of viewers.

8. The method of claim 7, wherein said information management system comprises a content delivery system that includes a storage system; said content delivery system being coupled to said network; and said storage system including said I/O resources and having at least one storage device or at least one partitioned group of storage devices.

9. The method of claim 8, wherein said managing comprises setting a cycle time of said storage device or partitioned group of storage devices to be greater than or equal to the service time of said storage device or partitioned group of storage devices.

10. The method of claim 9, wherein said managing further comprises setting a cycle time of said storage device or partitioned group of storage devices to maximize the number of simultaneous viewers of said continuous media data that is supported by said information management system.

11. The method of claim 8, wherein said method further comprises allocating at least one of said I/O resources between background system I/O activities and delivery of said continuous media data.

12. The method of claim 11, wherein said at least one of said allocated I/O resources comprise at least one of I/O capacity, buffer space, or a combination thereof.

13. The method of claim 12, wherein said background system I/O activities comprise at least one of said large file write operations, small file access operations, or a combination thereof.

14. The method of claim 13, wherein said method further comprises allocating said at least one I/O resources between said background system I/O activities and delivery of said continuous media data based at least in part on a variable resource parameter.

15. The method of claim 13, wherein said method further comprises dynamically allocating said at least one I/O resources between background system I/O activities and continuous media data delivery based at least in part on monitored background system processing activity.

16. The method of claim 15, wherein said method further comprises dynamically allocating said at least one I/O resources between background system I/O activities and continuous media data delivery in a manner that avoids interruptions in delivery of said continuous media data to existing viewers.

17. The method of claim 15, wherein said method further comprises selecting and terminating delivery of at least one stream of continuous media data to accommodate at least one background system I/O activity; wherein said at least one stream of continuous media data is selected for termination based at least in part on at least one priority-indicative parameter.

18. The method of claim 8, wherein said managing comprises at least one of performing I/O admission control, determining read-ahead size, or a combination thereof.

19. The method of claim 18, wherein said managing comprises performing said I/O admission control by monitoring the number of existing viewers served from said at least one storage device or partitioned group of storage devices and monitoring the data consumption rate of said existing viewers; balancing said I/O capacity with said buffer memory space based at least in part on said monitored number of existing viewers and said monitored data consumption rates of said existing viewers; and determining whether or not a capacity of said system is sufficient to support at least one additional viewer based at least in part on said balancing of said I/O capacity with said buffer memory space.

20. The method of claim 18, wherein said managing comprises determining said read-ahead size by monitoring the number of existing viewers served from said at least one storage device or partitioned group of storage devices and monitoring the data consumption rate of said existing viewers; balancing said I/O capacity with said buffer memory space based at least in part on said monitored number of existing viewers and said monitored data consumption rates of said existing viewers; setting a cycle time based at least in part on said balancing of said I/O capacity with said buffer memory space; and determining a number of read ahead data blocks based at least in part on said cycle time, said monitored data consumption rate, and a size of said data blocks.

21. The method of claim 18, wherein said managing comprises performing said I/O admission control by monitoring the number of existing viewers served from said at least one storage device or partitioned group of storage devices and monitoring the data consumption rate of said existing viewers; balancing said I/O capacity with said buffer memory space based at least in part on said monitored number of existing viewers and said monitored data consumption rates of said existing viewers; and determining whether or not a capacity of said system is

sufficient to support at least one additional viewer based at least in part on said balancing of said I/O capacity with said buffer memory space; and

wherein said managing further comprises determining said read-ahead size by monitoring the number of existing viewers served from said at least one storage device or partitioned group of storage devices and monitoring the data consumption rate of said existing viewers; balancing said I/O capacity with said buffer memory space based at least in part on said monitored number of existing viewers and said monitored data consumption rates of said existing viewers; setting a cycle time based at least in part on said balancing of said I/O capacity with said buffer memory space; and determining a number of read ahead data blocks based at least in part on said cycle time, said monitored data consumption rate, and a size of said data blocks.

22. The method of claim 6, wherein said modeling utilization of at least one of said I/O resources is performed using an analytical-based resource model.

23. The method of claim 6, wherein said modeling utilization of at least one of said I/O resources is performed using a measurement-based resource model.

24. The method of claim 8, wherein said at least one storage device comprises a disk drive or a partitioned group of disk drives; and wherein said modeling comprises modeling utilization of at least one of said I/O resources based at least in part on at least one system I/O performance characteristics associated with said I/O resources, said I/O performance characteristics comprising at least one of seek and rotation latency, estimated transfer rate, or a combination thereof.

25. The method of claim 3, wherein said method further comprises validating an estimated value of at least one of said system I/O performance characteristics by measuring a value of said

at least one system I/O performance characteristic, and comparing the measured value of said at least one system I/O performance characteristic to the estimated value of said at least one system I/O performance characteristic.

5

26. The method of claim 25, wherein said method further comprises reporting an alarm based at least in part on said comparison of the measured value of said at least one system I/O performance characteristic to the estimated value of said at least one system I/O performance characteristic.

10

27. A method of managing I/O resources for delivery of continuous media data to a plurality of viewers from a storage system including at least one storage device or at least one partitioned group of storage devices, said method comprising modeling utilization of at least one of said I/O resources; and managing at least one of said I/O resources based at least in part on said modeled utilization.

15

28. The method of claim 27, wherein said storage system comprises a part of a content delivery system coupled to a network; wherein said continuous media data is delivered to said plurality of viewers across said network; and wherein said storage system includes at least two storage devices or at least two partitioned groups of storage devices for delivery of said continuous media data.

20

25 29. The method of claim 28, wherein said content delivery system is coupled to said network at an endpoint of said network.

30. The method of claim 28, wherein said I/O resources comprise I/O capacity and buffer memory space of said information management system; wherein said method further comprises modeling said I/O capacity based at least in part on a workload distribution across said at least

30

two storage devices or at least two partitioned groups of storage devices; and wherein said managing comprises balancing said I/O capacity with said buffer memory space to ensure uninterrupted delivery of said continuous media data to said plurality of viewers from said at least two storage devices or said at least two partitioned groups of storage devices.

5

31. The method of claim 30, further comprising monitoring said workload distribution across said at least two storage devices or at least two partitioned groups of storage devices.

10

32. The method of claim 30, wherein said managing comprises ensuring that an amount of data consumed by each of said viewers during a cycle time of said storage system is less than or equal to the product of number of data blocks fetched from said at least two storage devices or partitioned groups of storage devices during a service time of said at least two storage devices or partitioned groups of storage devices, and a size of said fetched data blocks.

33. The method of claim 30, wherein said managing comprises ensuring that total available buffer space of said storage system is greater than or equal to the product of number of data blocks fetched from said at least two storage devices or partitioned groups of storage devices during a service time of said at least two storage devices or partitioned groups of storage devices, and a size of said fetched data blocks.

34. The method of claim 30, wherein said managing comprises ensuring that an amount of data consumed by each of said viewers during a cycle time of said storage system is equal to the product of number of data blocks fetched from said at least two storage devices or partitioned groups of storage devices during a service time of said at least two storage devices or partitioned groups of storage devices, and a size of said fetched data blocks; and by ensuring that total available buffer space of said storage system is greater than or equal to the product of number of data blocks fetched from said at least two storage devices or partitioned groups of

storage devices during a service time of said at least two storage devices or groups of storage devices, and a size of said fetched data blocks.

5 35. The method of claim 30, wherein said buffer memory space comprises a part of an integrated cache/buffer memory of said storage system.

10 36. The method of claim 35, further comprising monitoring a number of viewers that are reading data from said at least two storage devices or partitioned groups of storage devices out of the total number of viewers being served by said storage system; and wherein said managing comprises balancing said I/O capacity with said buffer memory space to ensure uninterrupted delivery of said continuous media data to said viewers reading data from said at least two storage devices or partitioned groups of storage devices.

15 37. The method of claim 35, wherein said method further comprises limiting a size of said buffer memory space by implementing a read-ahead buffer limit.

20 38. The method of claim 30, wherein said managing comprises setting a cycle time of said at least two storage devices or partitioned groups of storage devices to be greater than or equal to the maximal aggregate service time of said at least two storage devices or partitioned groups of storage devices.

25 39. The method of claim 38, wherein said managing further comprises setting a cycle time of said at least two storage devices or partitioned groups of storage devices to maximize the number of simultaneous viewers of said continuous media data that is supported by said information management system.

30

40. The method of claim 30, wherein said method further comprises allocating said I/O resources between background processing activities and delivery of said continuous media data.

5

41. The method of claim 30, wherein said managing comprises at least one of performing I/O admission control, determining read-ahead size, or a combination thereof.

10 42. The method of claim 41, wherein said managing comprises performing said I/O admission control by monitoring the number of existing viewers served from said at least two storage devices or partitioned groups of storage devices, and monitoring the data consumption rate of said existing viewers; balancing said I/O capacity with said buffer memory space based at least in part on said monitored number of existing viewers and said monitored data consumption rates of said existing viewers; and determining whether or not a capacity of said system is sufficient to support at least one additional viewer based at least in part on said balancing of said I/O capacity with said buffer memory space.

15 43. The method of claim 41, wherein said managing comprises determining said read-ahead size by monitoring the number of existing viewers served from said at least two storage devices or partitioned groups of storage devices, and monitoring the data consumption rate of said existing viewers; balancing said I/O capacity with said buffer memory space based at least in part on said monitored number of existing viewers and said monitored data consumption rates of said existing viewers; setting a cycle time based at least in part on said balancing of said I/O capacity with said buffer memory space; and determining a number of read ahead data blocks based at least in part on said cycle time, said monitored data consumption rate, and a size of said data blocks.

30

44. The method of claim 41, wherein said managing comprises performing said I/O admission control by monitoring the number of existing viewers served from said at least two storage devices or partitioned groups of storage devices, and monitoring the data consumption rate of said existing viewers; balancing said I/O capacity with said buffer memory space based at least in part on said monitored number of existing viewers and said monitored data consumption rates of said existing viewers; and determining whether or not a capacity of said system is sufficient to support at least one additional viewer based at least in part on said balancing of said I/O capacity with said buffer memory space; and

wherein said managing further comprises determining said read-ahead size by monitoring the number of existing viewers served from said at least two storage devices or partitioned groups of storage devices, and monitoring the data consumption rate of said existing viewers; balancing said I/O capacity with said buffer memory space based at least in part on said monitored number of existing viewers and said monitored data consumption rates of said existing viewers; setting a cycle time based at least in part on said balancing of said I/O capacity with said buffer memory space; and determining a number of read ahead data blocks based at least in part on said cycle time, said monitored data consumption rate, and a size of said data blocks.

45. The method of claim 30, wherein individual storage devices of said at least two storage devices or partitioned groups of storage devices comprise storage disk drives; and wherein said modeling comprises modeling utilization of at least one of said I/O resources based at least in part on at least one system I/O performance characteristics associated with said I/O resources, said I/O system performance characteristics comprising at least one of seek and rotation latency, estimated transfer rate, or a combination thereof.

46. The method of claim 28, wherein said modeling comprises modeling utilization of at least one of said I/O resources based at least in part on at least one estimated system I/O performance characteristics; and wherein said method further comprises validating an estimated value of at

51. The method of claim 50, wherein said information management system comprises a content delivery system that includes a storage system; said content delivery system being coupled to said network; and said storage system including said I/O resources and having at least one storage device or at least one partitioned group of storage devices.

52. The method of claim 51, wherein said performing admission control comprises using said modeled utilization of said I/O resources to determine whether or not said storage system has sufficient I/O capacity and sufficient buffer memory space to support a request for delivery of said continuous media data to a new viewer or an existing viewer returning from a cache state to an I/O state of said storage system, and admitting said viewer to said I/O state of said storage system if said sufficient I/O capacity and said sufficient buffer memory space exist; and wherein said determining read-ahead size comprises re-determining a read-ahead size for each of the existing viewers in said I/O state of said storage system if said new or returning existing viewer is admitted to said I/O state of said storage system, wherein said read-ahead size is determined based at least in part on I/O capacity of said storage system, buffer memory space of said storage system, the total number of existing viewers in said I/O state of storage system, and the data consumption rate of said total number of viewers in said I/O state of said storage system.

53. The method of claim 52, wherein said read-ahead size is determined for each of said existing viewers in said I/O state of said storage system prior to admitting said new viewer or returning existing viewer to said I/O state of said storage system.

54. The method of claim 53, wherein said storage system comprises at least two storage devices or at least two partitioned groups of storage devices.

55. The method of claim 53, wherein said method further comprises monitoring workload and utilization of said at least one storage device or partitioned group of storage devices; and wherein said modeling utilization of said I/O resources is based at least in part on said monitored workload and utilization of said at least one storage device or partitioned group of storage
5 devices.

56. The method of claim 53, wherein said method further comprises allocating said I/O resources between background processing activities and delivery of said continuous media data; and if said storage system has insufficient I/O capacity or insufficient buffer memory space to support a request for delivery of said continuous media data to said new viewer or said existing viewer returning from cache state of said storage system, then re-allocating at least a portion of said I/O resources from said background processing activities to said delivery of said continuous media data and again performing said admission control and said determining read-ahead size based at least in part on said re-allocated I/O resources.

57. The method of claim 51, wherein said storage system has an existing cycle time and an existing read-ahead size, and wherein said method further comprises leaving said existing cycle time and said existing read-ahead size unchanged if no new viewer is admitted to I/O state and no existing viewer is returned from cached state to I/O state; or performing the following steps if
20 an I/O request is received from a new viewer or existing viewer returning from said cached state of said storage system to said I/O state of said storage system:

modeling utilization of at least one I/O resources of said storage system based at least in part on admittance of said new viewer or existing viewer returning from said
25 cached state of said storage system to said I/O state of said storage system;

determining if a possible value or range of possible values of cycle time exist to balance said I/O capacity with said buffer memory space, wherein said I/O capacity is

balanced with said buffer memory space based at least in part on said modeled I/O resource utilization;

admitting said new viewer or existing viewer returning from said cached state of said storage system to said I/O state of said storage system if a possible value or range of possible values of cycle time exists to balance said I/O capacity with said buffer memory space;

refusing to admit said new viewer or existing viewer returning from said cached state of said storage system to said I/O state of said storage system if a possible value or range of possible values of cycle time is determined not to exist to balance said I/O capacity with said buffer memory space.

58. The method of claim 57, wherein said method further comprises refusing to admit said new viewer or existing viewer returning from said cached state of said storage system to said I/O state of said storage system if a predetermined range of possible values of cycle time is determined not to exist to balance said I/O capacity with said buffer memory space.

59. The method of claim 57, wherein if a possible value or range of possible values of cycle time exists to balance said I/O capacity with said buffer memory space, then performing one of the following steps prior to admitting said new viewer or existing viewer returning from said cached state of said storage system to said I/O state of said storage system:

leaving the existing value of cycle time and the existing value of read-ahead size unchanged if the existing value of cycle time is equal to said determined possible value of cycle time to balance said I/O capacity with said buffer memory space, or if the existing value of cycle time is within said determined range of possible values of cycle time to balance said I/O capacity with said buffer memory space; or

if the existing value of cycle time is not equal to said possible value of cycle time determined to balance said I/O capacity with said buffer memory space or is not within said range of possible values determined to balance said I/O capacity with said buffer memory space, then determining a new value of cycle time for said storage system that is equal to said determined possible value of cycle time to balance said I/O capacity with said buffer memory space, or that is within said determined range of possible values of cycle time to balance said I/O capacity with said buffer memory space; and determining a new value of read-ahead size for said storage system based at least in part on said new value of cycle time for said storage system.

60. The method of claim 57, wherein prior to refusing to admit said new viewer or existing viewer returning from said cached state of said storage system to said I/O state of said storage system, said method further comprises:

re-allocating at least a portion of said I/O resources from background processing activities to said delivery of said continuous media data;

modeling utilization of at least one I/O resources of said storage system based at least in part on said re-allocated I/O resources and based at least in part on admittance of said new viewer or existing viewer returning from said cached state of said storage system to said I/O state of said storage system;

determining if a possible value or range of possible values of cycle time exist to balance said I/O capacity with said buffer memory space, wherein said I/O capacity is balanced with said buffer memory space based at least in part on said modeled I/O resource utilization based at least in part on said re-allocated I/O resources; and then

admitting said new viewer or existing viewer returning from said cached state of said storage system to said I/O state of said storage system if a possible value or range of

possible values of cycle time based at least in part on said re-allocated I/O resources exists to balance said I/O capacity with said buffer memory space; or

refusing to admit said new viewer or existing viewer returning from said cached state of said storage system to said I/O state of said storage system if a possible value or range of possible values of cycle time based at least in part on said re-allocated I/O resources is determined not to exist to balance said I/O capacity with said buffer memory space;

61. The method of claim 60, wherein if a possible value or range of possible values of cycle time based at least in part on said re-allocated I/O resources exists to balance said I/O capacity with said buffer memory space, then performing one of the following steps prior to admitting said new viewer or existing viewer returning from said cached state of said storage system to said I/O state of said storage system:

leaving the existing value of cycle time and the existing value of read-ahead size unchanged if the existing value of cycle time is equal to said determined possible value of cycle time to balance said I/O capacity with said buffer memory space, or if the existing value of cycle time is within said determined range of possible values of cycle time to balance said I/O capacity with said buffer memory space; or

if the existing value of cycle time is not equal to said possible value of cycle time determined to balance said I/O capacity with said buffer memory space or is not within said range of possible values determined to balance said I/O capacity with said buffer memory space, then determining a new value of cycle time for said storage system that is equal to said determined possible value of cycle time to balance said I/O capacity with said buffer memory space, or that is within said determined range of possible values of cycle time to balance said I/O capacity with said buffer memory space; and determining a new value of read-ahead size for said storage system based at least in part on said new value of cycle time for said storage system.

62. The method of claim 57, wherein prior to refusing to admit said existing viewer returning from said cached state of said storage system to said I/O state of said storage system, said method further comprises:

re-allocating at least a portion of said I/O resources from background processing activities to said delivery of said continuous media data;

modeling utilization of at least one I/O resources of said storage system based at least in part on said re-allocated I/O resources and based at least in part on admittance of said existing viewer returning from said cached state of said storage system to said I/O state of said storage system;

determining if a possible value or range of possible values of cycle time exist to balance said I/O capacity with said buffer memory space, wherein said I/O capacity is balanced with said buffer memory space based at least in part on said modeled I/O resource utilization based at least in part on said re-allocated I/O resources; and then

admitting said existing viewer returning from said cached state of said storage system to said I/O state of said storage system if a possible value or range of possible values of cycle time based at least in part on said re-allocated I/O resources exists to balance said I/O capacity with said buffer memory space; or

refusing to admit said existing viewer returning from said cached state of said storage system to said I/O state of said storage system if a possible value or range of possible values of cycle time based at least in part on said re-allocated I/O resources is determined not to exist to balance said I/O capacity with said buffer memory space;

63. The method of claim 62, wherein if a possible value or range of possible values of cycle time based at least in part on said re-allocated I/O resources exists to balance said I/O capacity with said buffer memory space, then performing one of the following steps prior to admitting

said existing viewer returning from said cached state of said storage system to said I/O state of said storage system:

5 leaving the existing value of cycle time and the existing value of read-ahead size unchanged if the existing value of cycle time is equal to said determined possible value of cycle time to balance said I/O capacity with said buffer memory space, or if the existing value of cycle time is within said determined range of possible values of cycle time to balance said I/O capacity with said buffer memory space; or

10 if the existing value of cycle time is not equal to said possible value of cycle time determined to balance said I/O capacity with said buffer memory space or is not within said range of possible values determined to balance said I/O capacity with said buffer memory space, then determining a new value of cycle time for said storage system that is equal to said determined possible value of cycle time to balance said I/O capacity with said buffer memory space, or that is within said determined range of possible values of cycle time to balance said I/O capacity with said buffer memory space; and determining a new value of read-ahead size for said storage system based at least in part on said new value of cycle time for said storage system.

15
20 64. A method of modeling utilization of one or more I/O resources in an information delivery environment, comprising monitoring at least one of said system I/O performance characteristics associated with said I/O resources, and modeling utilization of at least one of said I/O resources based at least in part on said monitored I/O system performance characteristics.

25 65. The method of claim 64, wherein said information delivery environment comprises delivery of continuous media data to a plurality of viewers from an information management system; and wherein said I/O resources comprise I/O capacity and buffer memory space of said information management system.

66. The method of claim 65, wherein said information management system comprises a storage system, said storage system including said I/O resources and having at least one storage device or at least one partitioned group of storage devices.

5

67. The method of claim 66, wherein said information management system comprises a content delivery system coupled to a network; and wherein said information delivery environment comprises delivery of continuous media data across said network from said content delivery system to said plurality of viewers.

10

68. The method of claim 67, wherein said content delivery system comprises an endpoint content delivery system coupled to said network at an endpoint of said network.

15

69. The method of claim 66, wherein said storage system comprises at least two storage devices or two partitioned groups of storage devices; and wherein said at least one of said monitored system I/O performance characteristics comprise at least one of said system I/O performance characteristics at least partially reflective of workload distribution across said at least two storage devices or said at least two partitioned groups of storage devices.

20

70. The method of claim 69, wherein said at least one of said monitored system I/O performance characteristics comprise at least one of maximal aggregate consumption rate for each of said at least two storage devices or partitioned groups of storage devices, maximal aggregate number of viewers for each of said at least two storage devices or partitioned groups of storage devices, or a combination thereof.

25

71. The method of claim 70, wherein said method further comprises managing at least one of said I/O resources for delivery of said continuous media data to said plurality of viewers based at least in part on said modeled utilization.

5

72. The method of claim 71, wherein an actual value of workload Skew for at least one of said storage devices or at least one of said partitioned groups of storage devices is greater than or equal to about 2.

10

73. The method of claim 72, wherein said managing comprises balancing said I/O capacity with said buffer memory space to ensure uninterrupted delivery of said continuous media data to said plurality of viewers from said at least two storage devices or said at least two partitioned groups of storage devices.

15

74. The method of claim 73, wherein said managing comprises setting a cycle time of said two or more storage devices or partitioned groups of storage devices to be greater than or equal to the maximal aggregate service time of said two or more storage devices or partitioned groups of storage devices.

20

75. The method of claim 74, wherein said managing further comprises setting a cycle time of said two or more storage devices or partitioned groups of storage devices to maximize the number of simultaneous viewers of said continuous media data that is supported by said information management system.

25

76. The method of claim 71, wherein said method further comprises allocating said I/O resources between background processing activities and delivery of said continuous media data.

30

77. The method of claim 71, wherein said managing comprises at least one of performing I/O admission control, determining read-ahead size, or a combination thereof.

5

78. The method of claim 77, wherein said managing comprises performing said I/O admission control by monitoring the number of existing viewers served from said at least one storage device or group of storage devices, and monitoring the data consumption rate of said existing viewers; balancing said I/O capacity with said buffer memory space based at least in part on said monitored number of existing viewers and said monitored data consumption rates of said existing viewers; and determining whether or not a capacity of said system is sufficient to support at least one additional viewer based at least in part on said balancing of said I/O capacity with said buffer memory space.

10

15

79. The method of claim 77, wherein said managing comprises determining said read-ahead size by monitoring the number of existing viewers served from said at least one storage device or partitioned group of storage devices, and monitoring the data consumption rate of said existing viewers; balancing said I/O capacity with said buffer memory space based at least in part on said monitored number of existing viewers and said monitored data consumption rates of said existing viewers; setting a cycle time based at least in part on said balancing of said I/O capacity with said buffer memory space; and determining a number of read ahead data blocks based at least in part on said cycle time, said monitored data consumption rate, and a size of said data blocks.

20

25

80. The method of claim 77, wherein said managing comprises performing said I/O admission control by monitoring the number of existing viewers served from said at least one storage device or group of storage devices, and monitoring the data consumption rate of said existing viewers; balancing said I/O capacity with said buffer memory space based at least in part on said monitored number of existing viewers and said monitored data consumption rates of

30

said existing viewers; and determining whether or not a capacity of said system is sufficient to support at least one additional viewer based at least in part on said balancing of said I/O capacity with said buffer memory space; and

5 wherein said managing further comprises determining said read-ahead size by monitoring the number of existing viewers served from said at least one storage device or partitioned group of storage devices, and monitoring the data consumption rate of said existing viewers; balancing said I/O capacity with said buffer memory space based at least in part on said monitored number of existing viewers and said monitored data
10 consumption rates of said existing viewers; setting a cycle time based at least in part on said balancing of said I/O capacity with said buffer memory space; and determining a number of read ahead data blocks based at least in part on said cycle time, said monitored data consumption rate, and a size of said data blocks.

15 81. The method of claim 71, wherein individual storage devices of said at least two storage devices or partitioned groups of storage devices comprise storage disk drives; and wherein said at least one of said monitored I/O system performance characteristics comprise at least one of
20 seek and rotation latency, estimated transfer rate, or a combination thereof.

25 82. The method of claim 65, wherein said method further comprises validating an estimated value of at least one of said system I/O performance characteristics by comparing a monitored value of at least one system I/O performance characteristic to the estimated value of said at least one system I/O performance characteristic.

30 83. The method of claim 82, wherein said method further comprises reporting an alarm based at least in part on said comparison of the monitored value of said at least one system I/O

performance characteristic to the estimated value of said at least one system I/O performance characteristic.

5 84. A method of monitoring I/O resource utilization in an information delivery environment, comprising monitoring said I/O resource utilization at the logical volume level.

10 85. The method of claim 84, wherein said information delivery environment comprises delivery of continuous media data to a plurality of viewers from an information management system comprising a storage system, said storage system including said I/O resources and having at least one storage device or at least one partitioned group of storage devices.

15 86. The method of claim 85, wherein said monitoring of said I/O resource utilization comprises monitoring a workload of said at least one storage device or at least one partitioned group of storage devices at the logical volume level.

20 87. The method of claim 85, wherein said monitoring of said I/O resource utilization comprises monitoring system I/O performance characteristics of said at least one storage device or at least one partitioned group of storage devices at the logical volume level.

25 88. The method of claim 85, wherein said monitoring of said I/O resource utilization comprises constantly monitoring a workload of said at least one storage device or at least one partitioned group of storage devices at the logical volume level during run-time of said storage system; and wherein said method further comprises deciding to accept or reject at least one new I/O request based at least in part on said monitored workload.

30

89. The method of claim 85, wherein said method comprises monitoring at least one of maximal aggregate consumption rate for said at least one storage device or partitioned group of storage devices, maximal aggregate number of viewers for said at least one storage device or partitioned group of storage devices, or a combination thereof.

5

90. The method of claim 84, wherein said information delivery environment comprises delivery of continuous media data to a plurality of viewers from an information management system comprising a storage system, said storage system including said I/O resources and having at least two storage devices or at least two partitioned groups of storage devices.

10

91. The method of claim 90, wherein said information management system comprises a content delivery system coupled to a network; and wherein said continuous media data is delivered from said content delivery system to said plurality of viewers across said network.

15

92. The method of claim 91, wherein said content delivery system comprises an endpoint content delivery system coupled to said network at an endpoint of said network.

20

93. The method of claim 91, further comprising monitoring a workload distribution across said at least two storage devices or at least two partitioned groups of storage devices.

25

94. The method of claim 93, wherein said workload distribution is monitored by monitoring at least one of maximal aggregate consumption rate for each of said at least two storage devices or partitioned groups of storage devices, maximal aggregate number of viewers for each of said at least two storage devices or partitioned groups of storage devices, or a combination thereof.

30

95. The method of claim 94, wherein said I/O resources comprise I/O capacity; and wherein said method further comprises modeling said I/O capacity based at least in part on said monitored workload distribution across said at least two storage devices or partitioned groups of storage devices.

5

96. The method of claim 95, wherein said at least one of said I/O resources further comprise buffer memory space of said information management system; and wherein said method further comprises managing said I/O resources by balancing said I/O capacity with said buffer memory space to ensure uninterrupted delivery of said continuous media data to said plurality of viewers from said at least two storage devices or said at least two partitioned groups of storage devices.

10

97. The method of claim 96, wherein a maximum actual Skew value of at least one of said storage devices or partitioned groups of storage devices is greater than or equal to about 2.

15

98. The method of claim 96, wherein said buffer memory space comprises a part of an integrated cache/buffer memory of said storage system.

~

99. The method of claim 98, wherein said method further comprises limiting a size of said buffer memory space by implementing a read-ahead buffer limit.

20

100. The method of claim 93, wherein said I/O resources comprise I/O capacity and buffer memory space; and wherein said monitoring of workload distribution comprises monitoring a number of viewers that are reading data from each of said at least two storage devices or partitioned groups of storage devices out of the total number of viewers being served by said storage system; and managing said I/O resources by balancing said I/O capacity with said buffer memory space to ensure uninterrupted delivery of said continuous media data to said viewers reading data from said at least two storage devices or partitioned groups of storage devices.

25

30

09970457 100304
10030457 09970457

101. The method of claim 93, wherein said I/O resources comprise I/O capacity and buffer memory space; wherein said monitoring of workload distribution comprises monitoring the number of existing viewers served from each of said at least two storage devices or partitioned groups of storage devices, and monitoring the data consumption rate of said existing viewers; and wherein said method further comprises managing said I/O resources by balancing said I/O capacity with said buffer memory space based at least in part on said monitored number of existing viewers and said monitored data consumption rates of said existing viewers, and determining whether or not a capacity of said system is sufficient to support at least one additional viewer based at least in part on said balancing of said I/O capacity with said buffer memory space.

102. The method of claim 93, wherein said I/O resources comprise I/O capacity and buffer memory space; wherein said monitoring of workload distribution comprises monitoring the number of existing viewers served from each of said at least two storage devices or partitioned groups of storage devices, and monitoring the data consumption rate of said existing viewers; and wherein said method further comprises determining a read-ahead size by balancing said I/O capacity with said buffer memory space based at least in part on said monitored number of existing viewers and said monitored data consumption rates of said existing viewers, setting a cycle time based at least in part on said balancing of said I/O capacity with said buffer memory space; and determining a number of read ahead data blocks based at least in part on said cycle time, said monitored data consumption rate, and a size of said data blocks.

103. The method of claim 93, wherein said I/O resources comprise I/O capacity and buffer memory space; wherein said monitoring of workload distribution comprises monitoring the number of existing viewers served from each of said at least two storage devices or partitioned groups of storage devices, and monitoring the data consumption rate of said existing viewers; and wherein said method further comprises managing said I/O resources by balancing said I/O

capacity with said buffer memory space based at least in part on said monitored number of existing viewers and said monitored data consumption rates of said existing viewers, and determining whether or not a capacity of said system is sufficient to support at least one additional viewer based at least in part on said balancing of said I/O capacity with said buffer memory space; and wherein said method further comprises determining a read-ahead size by balancing said I/O capacity with said buffer memory space based at least in part on said monitored number of existing viewers and said monitored data consumption rates of said existing viewers, setting a cycle time based at least in part on said balancing of said I/O capacity with said buffer memory space; and determining a number of read ahead data blocks based at least in part on said cycle time, said monitored data consumption rate, and a size of said data blocks.

104. A method of monitoring I/O resource utilization for delivery of information to a plurality of viewers from an information management system including storage system I/O resources and at least one storage device or at least one partitioned group of storage devices; said method comprising logically monitoring workload of said at least one storage device or at least one partitioned group of storage devices.

105. The method of claim 104, wherein said logical monitoring comprises:

monitoring a number of viewers being served by at least one logical volume contained at least in part on said at least one storage device or partitioned group of storage devices, and monitoring the aggregated data consumption rates for said number of viewers being served by at least one logical volume contained at least in part on said at least one storage device or partitioned group of storage devices; and

determining an estimated total number of viewers for said at least one storage device or partitioned group of storage devices, and determining an estimated data consumption rate for

said estimated total number of viewers based at least in part on said monitored number of viewers and said monitored aggregated consumption rates for said viewers.

5 106. The method of claim 104, wherein said delivered information comprises continuous media data, and wherein said storage system includes two or more storage devices or two or more partitioned groups of storage devices for delivery of said continuous media data.

10 107. The method of claim 106, wherein said information management system comprises a content delivery system coupled to a network; and wherein said continuous media data is delivered from said content delivery system to said plurality of viewers across said network.

15 108. The method of claim 107, wherein said content delivery system comprises an endpoint content delivery system coupled to said network at an endpoint of said network.

109. The method of claim 107, wherein said logical monitoring comprises:

20 monitoring a number of viewers being served by at least a portion of each of said at least two storage devices or at least two partitioned groups of storage devices, and monitoring the aggregated data consumption rates for said number of viewers being served by said at least a portion of each of said at least two storage devices or at least two partitioned groups of storage devices;

25 monitoring a number of outstanding I/O requests for at least a portion of each of said at least two storage devices or at least two partitioned groups of storage devices;

determining an estimated total number of viewers for each of said at least two storage devices based at least in part on said number of viewers being served by at least a portion of each

of said at least two storage devices or at least two partitioned groups of storage devices, and said monitored number of outstanding I/O requests for at least a portion of each of said at least two storage devices or at least two partitioned groups of storage devices; and

5 determining an estimated aggregated data consumption rate for each of said at least two storage devices or at least two partitioned groups of storage devices based at least in part on said estimated aggregated data consumption rate for said number of viewers being served by said at least a portion of each of said at least two storage devices or at least two partitioned groups of storage devices, and said monitored number of outstanding I/O requests for at least a portion of
10 each of said at least two storage devices or at least two partitioned groups of storage devices.

110. The method of claim 109, further comprising:

15 determining an estimated workload distribution across said at least two storage devices or at least two partitioned groups of storage devices based at least in part on said monitored number of outstanding I/O requests for at least a portion of each of said at least two storage devices or at least two partitioned groups of storage devices;

20 wherein said estimated total number of viewers for each of said at least two storage devices or at least two partitioned groups of storage devices is determined based at least in part on said estimated total number of viewers being served by at least a portion of each of said at least two storage devices or at least two partitioned groups of storage devices, and said estimated workload distribution for each of said respective at
25 least two storage devices or at least two partitioned groups of storage devices; and

 wherein said estimated aggregated data consumption rate for each of said at least two storage devices or at least two partitioned groups of storage devices is determined based at least in part on said estimated aggregated data consumption rate for each of said
30 at least two storage devices or at least two partitioned groups of storage devices, and

estimated workload distribution for each of said respective at least two storage devices or at least two partitioned groups of storage devices.

5 111. The method of claim 107, wherein said logical monitoring comprises:

10 monitoring a number of viewers being served by each logical volume contained on said at least two storage devices or at least two partitioned groups of storage devices, monitoring the aggregated data consumption rates for said number of viewers being served by each logical volume contained on said at least two storage devices or at least two partitioned groups of storage devices, and monitoring the number of plex for each said logical volume on said at least two storage devices or at least two partitioned groups of storage devices;

15 monitoring a number of outstanding I/O requests for each said plex;

20 determining an estimated total number of viewers for each said plex based at least in part on said monitored number of plex for each logical volume and said monitored number of viewers for each logical volume;

determining an estimated aggregated data consumption rate for each said plex based at least in part on said monitored number of plex for each logical volume and said monitored aggregated data consumption rates;

25 determining an estimated total number of viewers for each of said at least two storage devices or at least two partitioned groups of storage devices based at least in part on said estimated total number of viewers for each said plex and said monitored number of outstanding I/O requests for each said plex; and

30 determining an estimated aggregated data consumption rate for each of said at least two storage devices or at least two partitioned groups of storage devices based at

least in part on said estimated aggregated data consumption rate for each said plex and said monitored number of outstanding I/O requests for each said plex.

5 112. The method of claim 111, further comprising:

determining a maximal number of outstanding I/O requests for each said plex, and
determining an estimated workload distribution across said at least two storage devices or
at least two partitioned groups of storage devices based at least in part on said monitored
10 maximal number of outstanding I/O requests for each said plex; and

wherein said estimated total number of viewers for each of said at least two
storage devices or at least two partitioned groups of storage devices is determined based
at least in part on said estimated total number of viewers for each said plex and said
15 estimated workload distribution for each of said respective at least two storage devices or
at least two partitioned groups of storage devices; and

wherein said estimated aggregated data consumption rate for each of said at least
two storage devices or at least two partitioned groups of storage devices is determined
based at least in part on said estimated aggregated data consumption rate for each said
20 plex and estimated workload distribution for each of said respective at least two storage
devices or at least two partitioned groups of storage devices.

25 113. The method of claim 111, wherein each of said storage devices comprise storage disk
drives.

114. The method of claim 107, further comprising determining a maximal total number of viewers per storage device and a maximal aggregated consumption rate storage device or per partitioned group of storage devices.

5

115. The method of claim 114, wherein said I/O resources comprise I/O capacity; and wherein said method further comprises modeling said I/O capacity based at least in part on said determined maximal total number of viewers per storage device or per partitioned group of storage devices, and said determined maximal aggregated consumption rate per storage device.

10

116. The method of claim 115, wherein said at least one of said I/O resources further comprise buffer memory space of said information management system; and wherein said method further comprises managing said I/O resources by balancing said I/O capacity with said buffer memory space to ensure uninterrupted delivery of said continuous media data to said plurality of viewers from said at least two storage devices or said at least two partitioned groups of storage devices; wherein said balancing is based at least in part on said determined maximal total number of viewers per storage device or per partitioned group of storage devices, and said determined maximal aggregated consumption rate per storage device or per partitioned group of storage devices.

15

20

117. The method of claim 116, wherein method further comprises performing I/O admission control, determining read-ahead size, or a combination thereof; wherein said performance of I/O admission control and determination of read-ahead size are based at least in part on said determined maximal total number of viewers per storage device or per partitioned group of storage devices, and said determined maximal aggregated consumption rate per storage device or per partitioned group of storage devices.

25

30

118. The method of claim 117, wherein said method comprises performing I/O admission control by determining whether or not a capacity of said system is sufficient to support at least one additional viewer based at least in part on said balancing of said I/O capacity with said buffer memory space.

5

119. The method of claim 117, wherein said method comprises determining read-ahead size by setting a cycle time based at least in part on said balancing of said I/O capacity with said buffer memory space; and determining a number of read ahead data blocks based at least in part on said cycle time, determined maximal aggregated consumption rate per storage device or per partitioned group of storage devices, and a size of said data blocks.

10

120. The method of claim 117, wherein said method comprises performing said I/O admission control by determining whether or not a capacity of said system is sufficient to support at least one additional viewer based at least in part on said balancing of said I/O capacity with said buffer memory space; and wherein said method further comprises determining read-ahead size by setting a cycle time based at least in part on said balancing of said I/O capacity with said buffer memory space; and determining a number of read ahead data blocks based at least in part on said cycle time, determined maximal aggregated consumption rate per storage device or per partitioned group of storage devices, and a size of said data blocks.

15

20

121. The method of claim 107, wherein said logical monitoring comprises monitoring the following system I/O performance characteristics for each logical volume, for each plex within a logical volume, and for each storage device or partitioned group of storage devices within a plex: (A) total number of viewers, (B) aggregated data consumption rate, (C) current weight of workload on a storage device in a plex, and (D) number of outstanding I/O requests for each storage device or partitioned group of storage devices.

25

30

122. The method of claim 107, wherein said method further comprises determining a workload weight distribution for each of said storage devices or partitioned group of storage devices based at least in part on said monitored number of outstanding I/O requests for each storage device or partitioned group of storage devices.

5

123. An I/O resource management system capable of managing I/O resources in an information delivery environment, comprising:

10

an I/O resource model capable of modeling utilization of at least one of said I/O resources; and

15

an I/O resource manager in communication with said I/O resource model, said I/O resource manager being capable of managing at least one of said I/O resources based at least in part on said modeled utilization.

20

124. The system of claim 123, wherein said I/O resource model is capable of modeling utilization of at least one of said I/O resources based at least in part on at least one of said system I/O performance characteristics associated with said I/O resources.

25

125. The system of claim 124, wherein the value of at least one of said system I/O performance characteristics is estimated.

30

126. The system of claim 124, wherein said I/O resource management system further comprises at least one I/O resource monitor in communication with at least one of said I/O resource manager or said I/O resource model, said I/O resource monitor being capable of monitoring the value of at least one of said system I/O performance characteristics.

127. The system of claim 123, wherein said I/O resources comprise at least one of file system resources, storage system resources, or a combination thereof.

128. The system of claim 123, wherein said information delivery environment comprises delivery of continuous media data from an information management system in communication with said I/O resource management system; wherein said I/O resources comprise I/O capacity and buffer memory space of said information management system; and wherein said I/O resource manager is capable of balancing said I/O capacity with said buffer memory space to ensure uninterrupted delivery of said continuous media data.

129. The system of claim 128, wherein said information management system comprises a storage system, said storage system including said I/O resources and having at least one storage device or at least one partitioned group of storage devices.

130. The system of claim 129, wherein said information delivery environment comprises delivery of continuous media data from said information management system to a network; and wherein said information management system comprises a content delivery system configured to be coupled to a network.

131. The system of claim 130, wherein said content delivery system is configured to be coupled to a network at an endpoint of said network.

132. The system of claim 129, wherein said I/O resource manager is capable of allocating at least one of said I/O resources between background system I/O activities and delivery of said continuous media data.

133. The system of claim 132, wherein said I/O resource management system further comprises at least one I/O resource monitor in communication with at least one of said I/O resource manager or said I/O resource model, said I/O resource monitor being capable of monitoring background system processing activity; and wherein said I/O resource manager is capable of allocating said at least one of said I/O resources between background system I/O activities and delivery of said continuous media data based at least in part on said monitored background system processing activity.

134. The system of claim 129, wherein said I/O resource manager is capable of at least one of performing I/O admission control, determining read-ahead size, or a combination thereof.

135. The system of claim 128, wherein said I/O resource model comprises an analytical-based resource model.

136. The system of claim 128, wherein said I/O resource management system further comprises at least one I/O resource monitor in communication with said I/O resource model; and wherein said I/O resource model comprises a measurement-based resource model.

137. An I/O resource management system capable of managing I/O resources for delivery of continuous media data to a plurality of viewers from a storage system including at least one storage device or at least one partitioned group of storage devices, said system comprising:

an I/O resource monitor, said I/O resource monitor being capable of monitoring at least one of said system I/O performance characteristics associated with said I/O resources;

an I/O resource model in communication with said I/O resource monitor, said resource model being capable of modeling utilization of at least one of said I/O resources based at least in part on said at least one of said monitored system I/O performance characteristics; and

an I/O resource manager in communication with said I/O resource model, said I/O resource manager being capable of managing at least one of said I/O resources based at least in part on said modeled utilization.

138. The system of claim 137, wherein said storage system comprises part of a content delivery system configured to be coupled to a network.

139. The system of claim 138, wherein said content delivery system is configured to be coupled to said network at an endpoint of said network.

140. The system of claim 138, wherein said storage system includes at least two storage devices or at least two partitioned groups of storage devices for delivery of said continuous media data.

141. The system of claim 140, wherein said resource monitor is capable of monitoring a workload distribution across said at least two storage devices or at least two partitioned groups of storage devices; wherein said I/O resources comprise I/O capacity and buffer memory space of said information management system; and wherein said I/O resource model is capable of modeling said I/O capacity based at least in part on a workload distribution across said at least two storage devices or two or more partitioned groups of storage devices; and wherein said resource manager is capable of balancing said I/O capacity with said buffer memory space to

ensure uninterrupted delivery of said continuous media data to said plurality of viewers from said at least two storage devices or said at least two partitioned groups of storage devices.

142. The system of claim 141, wherein said buffer memory space comprises a part of an integrated cache/buffer memory of said storage system; wherein said I/O resource monitor is capable of monitoring a number of viewers that are reading data from said two or more storage devices or partitioned groups of storage devices out of the total number of viewers being served by said storage system; and wherein said I/O resource manager is capable of balancing said I/O capacity with said buffer memory space to ensure uninterrupted delivery of said continuous media data to said viewers reading data from said two or more storage devices or partitioned groups of storage devices.

143. The system of claim 141, wherein I/O resource manager is capable of allocating said I/O resources between background processing activities and delivery of said continuous media data.

144. The system of claim 141, wherein I/O resource manager is capable of at least one of performing I/O admission control, determining read-ahead size, or a combination thereof.

145. The system of claim 141, wherein individual storage devices of said at least two storage devices or partitioned groups of storage devices comprise storage disk drives; and wherein said I/O resource model is capable of modeling utilization of at least one of said I/O resources based at least in part on at least one of said monitored system I/O performance characteristics associated with said I/O resources, said I/O system performance characteristics comprising at least one of seek and rotation latency, estimated transfer rate, or a combination thereof.

146. The system of claim 138, wherein said I/O resources comprise I/O capacity and buffer memory space of said storage system.

147. The system of claim 146, wherein said storage system comprises at least two storage devices or two partitioned groups of storage devices; and wherein said at least one of said monitored system I/O performance characteristics comprise at least one of said system I/O performance characteristics at least partially reflective of workload distribution across said at least two storage devices or said at least two partitioned groups of storage devices.

148. The system of claim 147, wherein said at least one of said monitored system I/O performance characteristics comprise at least one of maximal aggregate consumption rate for each of said at least two storage devices or partitioned groups of storage devices, maximal aggregate number of viewers for each of said at least two storage devices or partitioned groups of storage devices, or a combination thereof.

149. The system of claim 148, wherein said resource manager is capable of managing at least one of said I/O resources for delivery of said continuous media data to said plurality of viewers based at least in part on said modeled utilization.

150. The system of claim 149, wherein said resource manager is capable of balancing said I/O capacity with said buffer memory space to ensure uninterrupted delivery of said continuous media data to said plurality of viewers from said at least two storage devices or said at least two partitioned groups of storage devices.

151. The system of claim 149, wherein said I/O resource manager is capable of at least one of performing I/O admission control, determining read-ahead size, or a combination thereof.

152. The system of claim 138, wherein said I/O resource manager is capable of monitoring said at least one of said system I/O performance characteristics at the logical volume level.

153. The system of claim 152, wherein said I/O resource monitor is capable of monitoring said system I/O performance characteristics of said at least one storage device or at least one partitioned group of storage devices at the logical volume level.

154. The system of claim 152, wherein said I/O resource monitor is capable of monitoring at least one of maximal aggregate consumption rate for said at least one storage device or partitioned group of storage devices, maximal aggregate number of viewers for said at least one storage device or partitioned group of storage devices, or a combination thereof.

155. The system of claim 152, wherein said storage system includes at least two storage devices or at least two partitioned groups of storage devices.

156. The system of claim 155, wherein said I/O resource monitor is capable of monitoring a workload distribution across said at least two storage devices or at least two partitioned groups of storage devices.

157. An information delivery storage system, said storage system comprising:

a storage management processing engine that includes an I/O resource manager, a logical volume manager, and a monitoring agent; said I/O resource manager, said logical volume manager, and said monitoring agent being in communication; and

at least one storage device or group of storage devices coupled to said storage management processing engine;

wherein said information delivery storage system comprises part of an information management system configured to be coupled to a network.

5

158. The system of claim 157, wherein said storage management processing engine comprises at least one of said processing modules that are capable of performing at least one of I/O resource monitoring, I/O resource modeling, I/O resource management, or a combination thereof.

10

159. The system of claim 158, wherein said I/O resource manager comprises a storage system workload monitor.

15

160. The system of claim 159, wherein said monitoring agent is capable of monitoring number of outstanding I/O requests in at least one storage device or group of storage devices; and wherein said storage system workload monitor is capable of monitoring a number of viewers being served by at least one logical volume contained at least in part on said at least one storage device or partitioned group of storage devices, and monitoring the aggregated data consumption rates for said number of viewers being served by at least one logical volume contained at least in part on said at least one storage device or partitioned group of storage devices.

20

161. The system of claim 160, wherein said information management system comprises a content delivery system; wherein delivered information comprises continuous media data; and wherein said storage system includes two or more storage devices or two or more partitioned groups of storage devices for delivery of said continuous media data.

30

162. The system of claim 161, wherein said monitoring agent is capable of monitoring a number of outstanding I/O requests for at least a portion of each of said at least two storage devices or at least two partitioned groups of storage devices; and wherein said storage system workload monitor is capable of:

5
monitoring a number of viewers being served by at least a portion of each of said at least two storage devices or at least two partitioned groups of storage devices, and monitoring the aggregated data consumption rates for said number of viewers being served by said at least a portion of each of said at least two storage devices or at least two
10 partitioned groups of storage devices;

determining an estimated total number of viewers for each of said at least two storage devices based at least in part on said number of viewers being served by at least a portion of each of said at least two storage devices or at least two partitioned groups of
15 storage devices, and said monitored number of outstanding I/O requests for at least a portion of each of said at least two storage devices or at least two partitioned groups of storage devices; and

determining an estimated aggregated data consumption rate for each of said at
20 least two storage devices or at least two partitioned groups of storage devices based at least in part on said estimated aggregated data consumption rate for said number of viewers being served by said at least a portion of each of said at least two storage devices or at least two partitioned groups of storage devices, and said monitored number of outstanding I/O requests for at least a portion of each of said at least two storage devices
25 or at least two partitioned groups of storage devices.

163. The system of claim 162, wherein said workload monitor is further capable of:

30 determining an estimated workload distribution across said at least two storage devices or at least two partitioned groups of storage devices based at least in part on said

monitored number of outstanding I/O requests for at least a portion of each of said at least two storage devices or at least two partitioned groups of storage devices;

wherein said estimated total number of viewers for each of said at least two storage devices or at least two partitioned groups of storage devices is determined based at least in part on said estimated total number of viewers being served by at least a portion of each of said at least two storage devices or at least two partitioned groups of storage devices, and said estimated workload distribution for each of said respective at least two storage devices or at least two partitioned groups of storage devices; and

wherein said estimated aggregated data consumption rate for each of said at least two storage devices or at least two partitioned groups of storage devices is determined based at least in part on said estimated aggregated data consumption rate for each of said at least two storage devices or at least two partitioned groups of storage devices, and estimated workload distribution for each of said respective at least two storage devices or at least two partitioned groups of storage devices.

164. The system of claim 161, wherein said I/O resource manager has knowledge of the number of plex for each logical volume contained on said at least two storage devices or at least two partitioned groups of storage devices; wherein said monitoring agent is capable of monitoring a number of outstanding I/O requests for each said plex; and wherein said storage system workload monitor is capable of:

monitoring a number of viewers being served by each logical volume contained on said at least two storage devices or at least two partitioned groups of storage devices, monitoring the aggregated data consumption rates for said number of viewers being served by each logical volume contained on said at least two storage devices or at least two partitioned groups of storage devices;

determining an estimated total number of viewers for each said plex based at least in part on said monitored number of plex for each logical volume and said monitored number of viewers for each logical volume;

5 determining an estimated aggregated data consumption rate for each said plex based at least in part on said monitored number of plex for each logical volume and said monitored aggregated data consumption rates;

10 determining an estimated total number of viewers for each of said at least two storage devices or at least two partitioned groups of storage devices based at least in part on said estimated total number of viewers for each said plex and said monitored number of outstanding I/O requests for each said plex; and

15 determining an estimated aggregated data consumption rate for each of said at least two storage devices or at least two partitioned groups of storage devices based at least in part on said estimated aggregated data consumption rate for each said plex and said monitored number of outstanding I/O requests for each said plex.

20 165. The system of claim 164, wherein said monitoring agent is further capable of determining a maximal number of outstanding I/O requests for each said plex; and wherein said storage system workload monitor is capable of:

25 determining an estimated workload distribution across said at least two storage devices or at least two partitioned groups of storage devices based at least in part on said monitored maximal number of outstanding I/O requests for each said plex;

30 determining an estimated total number of viewers for each of said at least two storage devices or at least two partitioned groups of storage devices based at least in part on said estimated total number of viewers for each said plex and said estimated workload

distribution for each of said respective at least two storage devices or at least two partitioned groups of storage devices; and

determining an estimated aggregated data consumption rate for each of said at least two storage devices or at least two partitioned groups of storage devices based at least in part on said estimated aggregated data consumption rate for each said plex and estimated workload distribution for each of said respective at least two storage devices or at least two partitioned groups of storage devices.

166. The system of claim 164, wherein each of said storage devices comprise storage disk drives.

167. The system of claim 162, wherein said storage system workload monitor is capable of determining a maximal total number of viewers per storage device and a maximal aggregated data consumption rate storage device or per partitioned group of storage devices.

168. The system of claim 167, wherein said I/O resources comprise I/O capacity; and wherein said I/O resource manager is capable of modeling said I/O capacity based at least in part on said determined maximal total number of viewers per storage device or per partitioned group of storage devices, and said determined maximal aggregated consumption rate per storage device.

169. The system of claim 168, wherein said at least one of said I/O resources further comprise buffer memory space of said information management system; and wherein said I/O resource manager is capable of managing said I/O resources by balancing said I/O capacity with said buffer memory space to ensure uninterrupted delivery of said continuous media data to said plurality of viewers from said at least two storage devices or said at least two partitioned groups

of storage devices; wherein said balancing is based at least in part on said determined maximal total number of viewers per storage device or per partitioned group of storage devices, and said determined maximal aggregated consumption rate per storage device or per partitioned group of storage devices.

5

170. The system of claim 169, wherein said I/O resource manager is capable of performing I/O admission control, determining read-ahead size, or a combination thereof; wherein said performance of I/O admission control and determination of read-ahead size are based at least in part on said determined maximal total number of viewers per storage device or per partitioned group of storage devices, and said determined maximal aggregated consumption rate per storage device or per partitioned group of storage devices.

10

15

171. The system of claim 170, wherein said I/O resource manager is capable of performing I/O admission control by determining whether or not a capacity of said system is sufficient to support at least one additional viewer based at least in part on said balancing of said I/O capacity with said buffer memory space.

20

172. The system of claim 170, wherein said resource manager is capable of determining read-ahead size by setting a cycle time based at least in part on said balancing of said I/O capacity with said buffer memory space; and determining a number of read ahead data blocks based at least in part on said cycle time, determined maximal aggregated consumption rate per storage device or per partitioned group of storage devices, and a size of said data blocks.

25

173. The system of claim 170, wherein said resource manager is capable of performing said I/O admission control by determining whether or not a capacity of said system is sufficient to support at least one additional viewer based at least in part on said balancing of said I/O capacity with said buffer memory space; and wherein said method further comprises determining read-

30

ahead size by setting a cycle time based at least in part on said balancing of said I/O capacity with said buffer memory space; and determining a number of read ahead data blocks based at least in part on said cycle time, determined maximal aggregated consumption rate per storage device or per partitioned group of storage devices, and a size of said data blocks.

5

174. The system of claim 161, wherein said storage system workload monitor is capable of monitoring the following system I/O performance characteristics for each logical volume, for each plex within a logical volume, and for each storage device or partitioned group of storage devices within a plex: (1) total number of viewers, (2) aggregated data consumption rate, (3) current weight of workload on a storage device in a plex, and (4) number of outstanding I/O requests for each storage device or partitioned group of storage devices.

10

175. The system of claim 161, wherein said storage system workload monitor is capable of determining a workload weight distribution for each of said storage devices or partitioned group of storage devices based at least in part on said monitored number of outstanding I/O requests for each storage device or partitioned group of storage devices.

15

74

20